



**Cambridge
Elements**

Philosophy of Law

The Philosophy of Legal Proof

Lewis Ross

Contents

Introduction	1
1 Standards of Proof	1
2 Proof: Fixed or Flexible?	18
3 Should Proof Be Binary?	27
4 Legal Probabilism and Anti-Probabilism	36
5 Who Should Decide?	52
References	67

Introduction

Criminal courts can remove the liberty and even the lives of those accused of wrongdoing. Civil courts can deport asylum seekers, render companies employing thousands of people bankrupt, and remove children from the care of their parents. Using the right standards when deciding legal cases is therefore of the utmost importance in making sure those affected by trials receive a fair deal.

While legal standards can at first sound like a rather dry or technical subject, the project of deciding on the right standards raises fascinating and deep philosophical questions. These questions cut to the heart of debates about ethics, politics, psychology, and epistemology. Questions about legal standards force us to examine where the state's duty to protect society conflicts with the interests of those accused of wrongdoing. Understanding these debates also reveals how citizens can limit or control legal institutions that wield considerable – and potentially oppressive – power over them. Thinking about legal standards of proof is both a topic of great theoretical interest and one that affects the lives of everyone in society. The standards we use ultimately determine when the state can take away our freedom, our children, our property, and even our lives.

This Element is an introduction to the philosophy of legal proof. It aims to be accessible to students of both law and philosophy, presupposing no technical background in either subject.

The Element is organised around five questions.

- **Section 1** introduces the standards of proof and asks what justifies them.
- **Section 2** asks whether we should use different standards in different cases.
- **Section 3** discusses whether criminal trials should end in binary outcomes – guilty versus not guilty – or whether we should use more fine-grained verdicts.
- **Section 4** asks whether proof is simply about showing that something is probable or likely, concentrating on the famous ‘Proof Paradox’.
- **Section 5** considers who should be trusted with deciding the outcome of trials, focusing on the debate surrounding the jury system.

1 Standards of Proof

Societies use trials to resolve disagreements. These disagreements can be between private individuals, between corporations, and with the state itself. The disagreements can concern any number of issues, from mundane questions about who started a drunken fist fight, to disputes about the arcana of shipping law, questions about the results of an election, or criminal responsibility for murder or rape. Some disagreements end with trivial resolutions like trimming

a garden hedge found to be encroaching a neighbour's property. Other disputes end with outcomes of great severity, including the bankruptcy of corporations that employ thousands of people, the overturning of elections, or even the imposition of the death penalty.

Courts must be decisive when resolving disagreements. They must form a view on what happened and deliver a judgement about what should happen next. A court cannot end its work with a 'maybe' and a judge cannot throw up their hands and say they cannot decide one way or the other. Those accused of crimes must be punished or released; election results must stand or be overturned; at-risk children must be removed from the family home or kept where they are. This burden of deciding is difficult because courts are usually confronted with ambiguity. Trial participants often fundamentally disagree – this is why there is a trial! – and point to seemingly contradictory evidence.

Since courts must decide one way or another, they need some method for dealing with evidence and opinions that point in different directions. Courts therefore rely on various rules concerning when something should be taken as 'proven'. While questions about legal proof may sound dry or technical, they are in fact of central importance in legal and political philosophy. Think about it this way: the rules we choose to govern trials are really rules about when the state should use its monopoly of power to force people to do things – to go to prison, to surrender their children, to give up their assets. This section focuses on the most important rules used to decide legal trials – the *standards of proof*.

1.1 The Criminal Standard

A standard of proof is a rule used to determine when the evidence is strong enough for a positive judgement (e.g. finding someone guilty of a crime) to be appropriate. There are different standards of proof. One of the most important distinctions is between the standard used in the criminal law and in the civil law. Criminal proof is the primary focus of this Element, but we will discuss civil proof as we go along. All of this is easiest to appreciate at the level of concrete detail, so let's jump in.

Criminal law ranges over conduct that has been criminalised – such as murder, theft, assault, sexual offences, fraud, and so on. Criminal conduct is usually prosecuted *by* the state (rather than the victim) *against* an individual. Criminal law is distinctive because those judged guilty are open to receive the most serious sanctions available to the legal system. These sanctions are *punitive* and can involve the imposition of serious harms on the offender – such as imprisonment or (in some jurisdictions) corporal and capital punishment.

In a criminal trial, the standard of proof used in many jurisdictions is:

Criminal standard of proof = prove guilt beyond reasonable doubt¹

This standard instructs the court to convict the accused of a crime *only if* the evidence supports the guilt of the accused beyond a reasonable doubt. If this standard is not met, the accused must be judged not guilty (‘acquitted’). So, the criminal standard specifically governs ‘guilty’ verdicts, with ‘not guilty’ verdicts being returned whenever the standard of proof for guilt is not met.

Beyond reasonable doubt (‘BRD’ for short) is obviously a demanding standard. There are many things supported by good evidence that are nevertheless reasonable to doubt. For example, there may be good reasons to trust your sometimes unreliable friend when they promise to meet you at 7 p.m. for a beer. But it might also be reasonable to harbour doubts. The BRD standard tells us to convict only if there are no reasonable doubts. This means that even if there is some evidence that the accused is guilty – even if you think there is a ‘good chance’ they are guilty – the court should release the accused so long as there is reasonable doubt. There are much less demanding standards of proof we might use. Indeed, there are other legal standards of proof used outside the trial context. For example, the standard used in different jurisdictions within the United Kingdom to decide whether the police can stop and search somebody is ‘reasonable grounds for suspicion’. Clearly, it would be a rather different criminal justice system if criminal courts imprisoned anybody against whom there was reasonable suspicion! Beyond reasonable doubt is a demanding standard of proof.

1.1.1 The *Actus Reus* and *Mens Rea*

Since it will be important later, it is worth saying more about *what* must be proven against the BRD standard to establish criminal guilt. In the ‘common law’ legal systems we focus on, there are two components jointly required to prove someone has committed a crime.² These two components are known by their Latin names – the *actus reus* and the *mens rea*. While the Latin may be unhelpful, the basic idea is pretty straightforward.

First, the *actus reus* is the ‘active’ part of the crime – the action or conduct that the law prohibits. Here are some rough examples. For theft, the prohibited action is taking the possessions of another without authority, for rape the prohibited action is non-consensual sexual intercourse, for murder the

¹ I focus on Anglo-American systems, which have served as the model for many international institutions (like the International Criminal Court). Other jurisdictions use different phrasings for the criminal standard, but generally have a similarly demanding approach.

² Common law legal systems are characteristically defined by reliance on what judges have said in previous cases – ‘precedent’ – to interpret and create law.

prohibited action is causing the death of another person, and so on. Perhaps obviously, criminal conviction requires proving that the accused performed the *actus reus* – the action (or sometimes omission) that constitutes a crime.³

Second, the *mens rea* is the ‘mental’ component of the crime. For example, ‘intent’ is a common *mens rea* found in the definition of many crimes. For a crime with a *mens rea* of intent, the *actus reus* has to be performed intentionally. This demonstrates where the *actus reus* and the *mens rea* can diverge; for example, someone can cause a death unintentionally. Something is only a crime when there is the right kind of unity between *actus reus* and *mens rea*, between the action that is performed and the mental state underpinning it. There are other mental states beyond intent that can be the *mens rea* for various crimes and we will come back to these later, but this simple account should be enough to go on for now.

To *prove* a crime, the prosecution has the burden of proof to establish *both* the *actus reus* and the *mens rea*, against the standard of proof. So, take the example of theft. To prove the crime of theft you must show that it is beyond reasonable doubt that the accused took the property *and* that it is beyond reasonable doubt that they intended to do so.

1.2 The Civil Standard

Let us now turn to the second main standard of proof, the one used in the civil law.

The civil law regulates the wide variety of non-criminal disputes adjudicated by the legal system. This includes contractual disputes, employment law, corporate law, family law (e.g. disputes about divorce), disputes about ‘negligence’ (often called torts), and constitutional law. Civil cases can be pursued by almost any person or legal entity against almost any other person or legal entity. In most civil cases, the standard of proof is the following:

Civil standard of proof = prove your case on the *balance of probabilities*

A common way to think about the balance of probabilities is just that it means ‘more likely than not’. So, take an example from employment law. An employer should be held liable for breaching their obligations (e.g. failing to provide safety equipment) just so long as the court finds it more likely than not that they failed to provide safety equipment and had an obligation to provide it.

³ A criminal omission might occur where a duty to prevent something is imposed by law – for example, as might apply to public office holders. ‘Attempts’ can also be a criminal *actus reus*, such as attempting to kill somebody. There are tricky issues in determining when someone has committed a criminal attempt rather than just having a vague plan or intention.

The sanctions that result from losing a civil case are extremely varied. The most common is being compelled to pay compensation. Financial penalties can be vast and serious, leading to bankruptcy or impoverishment. But civil cases cannot generally lead to someone being imprisoned or subjected to other types of punitive treatment. A general rule of thumb in civil law is that compensation aims to provide ‘restitution’ rather than punishment; it aims to put the party that was harmed in as good a place as they would have been had you not harmed them. For example, the court might try to estimate how much your interests were set back by having your contract breached and ask the other side to make up for it in monetary terms.

Clearly, the criminal standard is harder to satisfy than the civil standard – you can reckon something is ‘more likely than not’ while still having reasonable doubts. I might think it is more likely than not that my sometimes unreliable friend will turn up for our beer, but due to their track record I may have reasonable doubts. It is entirely consistent for evidence to be strong enough to satisfy the civil standard but not the criminal standard. Indeed, there are instances where this happens. Sometimes people are found not guilty of sexual assault in a criminal court (and hence not subject to punitive treatment like imprisonment) but found liable for sexual assault in a civil court (and hence asked to pay monetary compensation).

1.3 Justifying the Standards

How do we come up with the different standards of proof? In truth, the standards of proof have been heavily influenced by historical circumstance and have evolved piecemeal over time. Especially in common law countries where judicial opinions in individual cases influence the way the law evolves – the system of ‘precedent’ – the history of legal rules is often convoluted rather than the product of a single design. (Of course, convoluted and complex evolutionary processes do not necessarily make for worse products.) Legal history has, for me, a compulsive nerdish attraction because it shows how fragile and often accidental the way that the law works is. As we will see throughout the Element, it is also a source of stories, where idiosyncratic characters find themselves in the courts and change the way that entire states have operated.

The history of the ‘beyond reasonable doubt’ standard is much discussed, and the introduction of this terminology happened gradually rather than all at once. The language of ‘reasonable doubt’ evolved from philosophical discussions throughout the 1600s–1800s that worried about the fact that it might be impossible to prove almost anything with absolute (or ‘metaphysical’) certainty. (Remember Descartes!) Instead, it was thought that proof of everyday matters,

where there is always a chance of error, should be linked to the conscience of the person looking at the evidence; this is sometimes called proving something with ‘moral’ certainty.⁴ From this religiously inflected language, the secular idea of proof beyond reasonable doubt emerged. The connection between proof and the individual conscience will reappear throughout the Element.

As philosophers we are not *primarily* interested in how legal standards came to be the way they are. Rather, we are interested in how they ought to be. We can break this up into two related questions:

- (1) Can we reconstruct a justification for the current standards of proof?
- (2) In light of how we justify the standards, are they set in the best way?

A natural way to try and justify the different standards of proof is to think about how bad different types of mistakes would be. This is because what standards of proof do, in effect, is strike a balance between different types of error. To see this, consider criminal law.

When criminal courts make decisions – either finding someone guilty or not guilty – they can get it wrong in two different ways. One mistake is *convicting an innocent* person and punishing them for a crime they did not commit. Another type of mistake is *acquitting a guilty* person, allowing a criminal to walk free unpunished. Both of these mistakes are bad, regardless of what type of moral theory you endorse. Convicting the innocent and acquitting the guilty (generally) has bad consequences; the former inflicts misery, while the latter can mean that a dangerous person is released back into the community. It also seems to be unfair even aside from the consequences; the innocent don’t *deserve* to be punished, while the guilty might *deserve* punishment. (Of course, this assumes a rather traditional view about the value of punishment. Some might wonder whether punishing the guilty really has good consequences or whether it really is true that those who commit crimes *deserve* to be harmed. While I have some sympathy with this outlook, scepticism about punishment will mostly wait for another day.)

1.4 Two Types of Mistake

Setting the standard of proof involves a balancing act. The harder we make convicting someone of a crime, the less often we will convict people for crimes they didn’t commit. Demanding standards provide protection to the accused. However, by making the standard harder to satisfy, we thereby also make it more likely we will acquit people of crimes they *are* guilty of. And, of course, the converse is equally true. The lower the standard, the easier it is to convict; we end up blaming

⁴ Shapiro 1986.

more people who are genuinely guilty, but also, we make it more likely that we will convict the innocent. How should we perform this balancing act?

One approach might be to suppose that different decisions have some type of expected value or benefit. Then, once we have thought about how large or small these values are, we try to set the standard in a way that would maximise the expected overall value of all the decisions that courts make.

That's quite abstract, so here's an analogy. Suppose you are a fisherman deciding what net to use. You are only after a certain type of fish, red snappers. They are the only fish you can sell at the market; catching other types of fish drains your resources and harms the fish unnecessarily. If you use a very fine net, you'll let fewer red snappers escape, even snappers that are small and difficult to ensnare. But you'll also catch other types of fish too, ones you don't want. If you use a coarser net, you avoid mistakenly landing the unwanted fish you would catch with the fine net, but you also let more precious red snappers get away. What type of net should you use? Well, it depends on the relative value of catching the snappers compared to the expense of catching the unwanted fish. If there's a big difference in value, then we might be justified in using a very coarse or a very fine net. In the criminal law, in effect, we currently use a very coarse net. We leave aside many finer nets – namely, weaker standards of proof – that would catch more guilty people. But why?

1.5 Blackstone's Asymmetry

The following idea by the English jurist, William Blackstone, is often used by way of justification:

Blackstone's asymmetry: It is much worse to mistakenly convict an innocent person than to mistakenly acquit a guilty person.

Blackstone himself suggested it would be *ten times* worse to convict an innocent than acquit the guilty, famously saying: 'All presumptive evidence of felony should be admitted cautiously, *for the law holds that it is better that ten guilty persons escape than that one innocent suffer.*'⁵ It's an interesting historical question why Blackstone's 10:1 ratio became the canonical version of the asymmetry known to all law students. In fact, throughout history we find a dizzying number of attempts to formulate a ratio. Even the Book of Genesis contains a passage in which Moses asks how many innocents would have to be present in Sodom in order to

⁵ Blackstone 1827, book four, chapter 27 (emphasis added).

prevent it from being destroyed by God.⁶ But something about the 10:1 ratio has a ring of plausibility. There are ways to build formal models designed to maximise the expected utility of our decisions – a branch of philosophy called ‘decision theory’ – which suggest that Blackstone’s 10:1 ratio recommends a 90 per cent level of confidence as the best level at which to set the threshold for conviction.⁷ Intriguingly, 90 per cent confidence in guilt is roughly where some people settle when attempting to quantify the BRD standard of proof. (We’ll have more on probabilistic approaches to proof later on.) This 90 per cent confidence level also matches up, roughly, with what some empirical surveys have said about how BRD is interpreted by judges.⁸

However, there is a big question facing Blackstone’s seductive asymmetry. How can we justify the claim that it is *much worse* to convict an innocent than mistakenly release the guilty?

One way to think about criminal justice is to focus on what people ‘deserve’ irrespective of the consequences (a view sometimes called ‘*retributivism*’). For instance, punishing an elderly criminal whose victims have long since died might be expensive and yield little obvious future-oriented benefit. Yet, some might think that seeking to convict such a person is the just thing to do regardless of whether it leads to any particular beneficial consequences. There are different ways to elaborate on this idea of criminal justice aiming to give people what they deserve. But, it isn’t immediately obvious that focusing on what people deserve justifies a very demanding standard of proof. While it is true that convicting the innocent fails to give people what they deserve, so does mistakenly acquitting the guilty. Consider the following. One worry, to which we will return repeatedly as a matter of policy interest, is the low conviction rate for sexual crimes. There is a striking drop-off rate in the number of sexual offences reported to the police against the number that are prosecuted in the courtroom. For example, in England and Wales, a recent report claims that less than 2 per cent of complaints lead to a conviction.⁹ One possibility is that the high standard of proof is partly responsible for the low conviction rates for certain types of crimes. Proof beyond a reasonable doubt can be difficult in the context of sexual criminality due to the often private nature of such crimes. If

⁶ See Volokh 1997 for an entertaining discussion.

⁷ On such approaches, see Kaplan 1968; Lillquist 2002. ⁸ For example, see Solan 1999.

⁹ HM Government 2021, 7. See Thomas 2023 for empirical work on the conviction rate for sexual offences at trial: her findings suggest that criminal trials themselves may not be site of the problem, with conviction rates above 65 per cent in recent years. Of course, the lesson of the low complaint-to-conviction ratio is that many allegations never make it to trial. The standards used in trials influence both police and prosecutorial decisions. For instance, prosecutors will decline to pursue a charge precisely because they believe it is unlikely to meet the standard of criminal proof used at trial.

you think that the ‘base rate’ of offenders – that is, the number of people who are actually guilty – is higher than the purported 2 per cent being convicted after a complaint, then you might worry that the current standard is preventing these people from ‘getting what they deserve’. So, focusing on what people deserve doesn’t seem to provide an obvious vindication of using a coarse net.

Another way we might think through this issue is in terms of harm. Does much more misery, pain, and unhappiness result from convicting an innocent person compared to letting a guilty person walk free? It’s very hard to answer this in general terms. All crimes are different, all victims are different, and all accused are different. You might be sceptical about our ability to make confident predictions about the amount of harm caused by different criminal justice errors. Still, in many areas of life we can sensibly rely on rough generalisations – it’s worse for a doctor to mistakenly amputate your leg than your little finger, and better to cure you of a terminal illness than of a cold. These generalisations can be used to drive policy even if they admit of idiosyncratic exceptions, such as the pianist who would prefer to keep the finger rather than the leg. Perhaps it is *generally* true that more harm results by locking an innocent person up – inflicting a great deal of fruitless misery on them – compared to letting an innocent person go free without them ‘getting what they deserve’ or being rehabilitated?

A good way to scrutinise a claim like this is to consider an argument against it. Some have worried that we might be setting the standard for conviction *too high* in the criminal law and that we are overlooking or minimising the harms that follow from mistakenly releasing the guilty. Larry Laudan – famous first for his work in the philosophy of science, before turning later in his career to legal philosophy – has used this concern to develop a provocative argument against the BRD standard.¹⁰

1.6 Consequences-Based Arguments against Beyond Reasonable Doubt

The state has a duty to protect society from harm. When the state releases a guilty person by mistake, it enables a risk of harm to the rest of society. The risk is that the guilty person will reoffend against members of the community. If the court had got it right and found the accused guilty, the person would have been incapacitated through imprisonment. Laudan thinks that the BRD standard skews too heavily in favour of protecting innocent people from the potential harm of false conviction, rather than protecting innocent people from the harms caused by mistakenly released criminals. Indeed, he canvasses criminological

¹⁰ The following reconstruction draws on Laudan 2003, 2006, 2011.

evidence – from the US – that purports to show that you are much more likely to fall victim to a violent recidivist than be mistakenly convicted of a crime. From the perspective of someone looking to minimise their own risk, you could use this observation to argue for *lowering* the standard of proof. After all, if the chance of being mistakenly locked up is tiny, but the chance of being harmed by a wrongdoer who could otherwise have been imprisoned is comparatively large, then weakening the standard might seem entirely sensible from the perspective of harm-minimisation.

Laudan's interpretation and use of statistics has been trenchantly criticised.¹¹ There is also a lot left out by the argument I sketched – for example, assumptions about the underlying ratio of guilty and innocent people, about how frequent different types of mistakes currently are, and about whether mistakes and benefits are distributed unevenly across society. Many of the supposed benefits of identifying the guilty are also controversial and uncertain. Punishment is the main culprit here. Punishment, we are told, has various benefits: it rehabilitates, it deters future wrongdoing by the accused who will wish to avoid repeat punishment, and it deters would-be wrongdoers by making a life of crime generally unattractive. But the empirical evidence on rehabilitation and deterrence, in many cases, does not support these benefits.¹² Often, the threat of harsh sentences seems to do little to reduce crime rates. The same anti-criminal benefit could often be achieved by increasing wages or employment rates. Moreover, prison can even have a criminogenic effect – making people more likely to reoffend rather than rehabilitating them. This is especially true in the many states across the world that have poorly maintained and under-resourced prisons.

Still, even if Laudan's argument rests on shaky or even false empirical claims, the philosophical point remains important. Things can change. If things *were* as Laudan describes them, would this be a compelling reason to lower the standard of criminal proof below 'beyond reasonable doubt'? For example, at one point Laudan suggests a criminal standard of around 65 per cent confidence might be appropriate, a standard not much stronger than the 'more likely than not' standard of civil law.

Laudan's argument is just one of various consequence-based arguments that we could use to criticise the high standard of criminal proof. In addition to the costs imposed by reoffenders, there are various other costs that arguably result from a high standard of proof. Trials cost money and time, as well as being hard on victims, so the legal system tries to encourage those accused to admit their guilt

¹¹ Gardiner 2017.

¹² For meta-analyses on the effects of punishment, see Paternoster 2010; Nagin 2013; Chalfin and McCrary 2017.

before a trial. (Many jurisdictions offer lesser punishments for an early guilty plea and others have systems of ‘plea bargaining’ where those who agree to plead guilty are charged with less serious offences.) The standard of proof affects the rationality of pleading guilty. The higher the standard of proof, the more rational it becomes for a guilty person to take a gamble and plead not guilty – to try and escape any punishment. Daniel Epps argues that the BRD standard encourages the guilty to ‘chance their luck’.¹³ If this is right, the overall effect of a high standard of proof is more time and money spent on needless trials, and more guilty people getting lucky and escaping justice. Others discuss the possibility of high standards of proof undermining public confidence in the legal system (because people think too many guilty persons are acquitted) and worry about overly high standards of proof demoralising those responsible for apprehending criminals or encouraging them to use improper methods to gather evidence.¹⁴

1.7 Defending Beyond Reasonable Doubt?

How can we respond to these arguments? A straightforward way to respond to consequences-based arguments is to fight fire with fire, claiming that the consequences would actually be *worse* if we lowered the standard of criminal proof. Perhaps a lower standard of proof would actually diminish trust in the criminal justice system, because people would perceive it as less accurate? This may be true, but it is hard to be certain that reducing public confidence would have worse consequences than releasing violent reoffenders in the way Laudan complains about. Of course, uncertainty itself may be an argument against making radical changes – if in doubt, it might be best to leave things as they are, given that we have a criminal justice system that functions to some degree. Still, this conservatism is not entirely satisfying.

Another approach might focus on the following fact: one difference between punishing the innocent and the harm caused by reoffending criminals is that the first is actively imposed by the state itself while the latter is just something the state fails to prevent. Perhaps the state is not ultimately *responsible* for harms caused by reoffenders (even though it could conceivably prevent them), while the state *is* responsible for the bad that results from mistakenly punishing the innocent.¹⁵ A large literature exists on the moral difference between ‘doing’ versus ‘allowing’.¹⁶ It is deeply unclear whether this distinction has any general moral significance, so it is therefore not clear whether it can be used to respond

¹³ Epps 2015. This goes both ways; the lower the standard of proof, the more rational it becomes for an innocent person to plead guilty.

¹⁴ Kitai 2003. Kitai herself isn’t ultimately convinced by this worry.

¹⁵ Kitai 2003 defends an argument along these lines.

¹⁶ For example, see Woollard and Howard-Snyder 2022.

to Laudan. Moreover, there are reasons for doubting the usefulness of an actions/omissions-type distinction here. As Laudan himself points out, the state actively *does* something when it creates and sustains a criminal justice system where accused persons are released even when there is reasonably strong evidence they may commit further crimes.¹⁷

A further point in favour of prioritising protecting the accused is that, in criminal trials, the prosecution has most of the power. The prosecution is backed up by the might of the state, the police, and a skilled cadre of lawyers – in almost all cases, the prosecution has greater resources than the accused. Moreover, as Richard Lippke points out, the defence is at a rhetorical disadvantage. As he puts it: ‘Defence attorneys are hired by the accused to represent them, so of course they must say that their clients are innocent. . . . Defence attorneys are apt to be seen as little more than hired guns.’¹⁸ Prosecutors, on the other hand, are in court – in theory – because the state thinks the evidence demonstrates the guilt of the accused. So, the prosecution might benefit from an automatic (and sometimes unearned) trust. An argument for prioritising the interests of the accused is to protect the community from abusive or incompetent exercise of state power. Requiring crime to be proved beyond reasonable doubt is a final protection against such malfeasance. This illustrates a general tension that often arises between two reasonable perspectives within criminal justice and philosophy of law generally. Seemingly compelling arguments that emphasise the importance of the state protecting us from harm (e.g. from criminality), often conflict with another important perspective, namely, the importance of individuals protecting themselves from a state that is too eager to use – and perhaps abuse – its stranglehold on policing and punishment.

This argument against Laudan is promising. But it isn’t clear that a very high standard of proof is the best way to protect against the misuse of state power. As I’ll discuss in [Section 5](#), the jury system might serve this function irrespective of what standard of proof is used.

It is good philosophical practice to test positions by asking what they would say in more extreme circumstances. So, another way to scrutinise consequences-based arguments for lowering the standard of proof is to consider what they would recommend if the empirical situation worsened. Suppose things were not only as dangerous as Laudan suggests but rather *more* dangerous. Presumably, there would come a point where, according to the logic of Laudan’s argument, the underlying empirical situation would not just recommend moderately weakening, but radically weakening, the standard of proof. For instance, recall the ‘reasonable suspicion’ standard regulating police

¹⁷ Laudan 2011, 222. ¹⁸ Lippke 2010, 478.

searches. On Laudan's premises, there may come a point where we would be better off – from the perspective of harm-reduction – endorsing a criminal standard of proof as weak as reasonable suspicion of guilt.

Imprisoning people if there is only a reasonable suspicion of guilt seems unacceptable. Not only does it seem like a recipe for state oppression, it also seems like an objectionable way of sacrificing the interests of an individual accused person for the greater good of the collective. Something in Laudan's perspective, I think, has gone awry.

It is true that we do not use the strongest imaginable safeguards against convicting the innocent. We could use larger juries, always requiring unanimity to convict. We could impose tighter restrictions on when incriminating evidence is admissible (e.g. requiring it be verified by independent sources). We could have automatic post-trial reviews of every conviction. All of these changes would make it harder to convict the innocent. The fact that we don't do these things recognises the fact that we must ensure there is a reasonable prospect of securing convictions. Yet, I think there is a limit in how far we can weaken the safeguards against convicting the innocent before crossing an important moral line.

1.8 Criminal Proof and Community Belief

Rejecting consequences-based arguments for a low standard of proof is different from saying that consequences don't matter for the criminal law at all. We can – and probably should – grant that consequences are generally important in criminal justice. For example, expected consequences might matter for how we should punish people after they are convicted (e.g. punish in a way most effective for rehabilitation) or in setting the rules for granting parole. But, I think that justice requires that we do not appeal to the expected consequences of punishing *until* we have fairly decided whether the person deserves to be punished in the first place.¹⁹ The consequences-based reasons only come into play *after finding the accused guilty* – not when we set the standard of proof used to determine guilt in the first place.

With this thought in mind, I defend a way to think about setting the standard of criminal proof that does not appeal to consequences. We might call the approach a 'blame signalling' view.

Courts play a fundamentally social role – they exist to settle disagreements on behalf of their community. This means that, ideally, the verdicts that courts reach should be an effective social signal that the disagreement has in fact been settled. If the justice system is fair and commands their trust, people in the

¹⁹ For a similar thought, see [Walen 2015](#), 427.

community should be able to take the content of a court's verdict as good evidence about the truth of the matter – even if they haven't had the chance to consider the evidence for themselves. (Of course, given the thousands of disagreements that need to be solved by courts, not every citizen *can* consider every case for themselves!) In other words, a court's verdict should ideally be a proxy for what an individual in the community would have believed if they had considered the matter for themselves.²⁰ To see the importance of this, consider what happens if this *doesn't* happen. If courts regularly found people guilty of crimes but the community did not then believe the convicted were guilty, courts would lose legitimacy. They would not fulfil the role of settling disagreements for the community in question – the guilt of those convicted in court would remain an open question in the mind of the community.

Building on this idea we can make sense of an important idea in philosophy of law. This is the idea that the standard of criminal proof has a close connection with the standards of rational *belief*.²¹

Not all standards of proof have an essential connection to belief. For example, think about the 'reasonable suspicion' standard used for stop and search. Clearly, you can reasonably suspect something is the case without believing it (e.g. you might reasonably suspect your date has stood you up but not quite believe it yet!). But proving that someone is guilty of a crime, I think, ought to be different. You should not find someone guilty without fully believing that they are guilty. But why?

One influential position in moral psychology draws a connection between belief and moral *blame*. For example, Lara Buchak has argued that one thing that is distinctive about belief is the role it plays in legitimating blame. Specifically, Buchak argues that one thing that sets belief apart from other attitudes is that we must believe that someone is responsible for something before blaming them.²² Let's think about this idea a bit more closely.

In many contexts, we proportion our reaction to the evidence along a spectrum. For example, if it's only 25 per cent likely to rain you act one way (risk shorts and T-shirt), if it's 50 per cent likely to rain you act differently (take an umbrella just in case), and if it's 95 per cent likely to rain you act differently yet again (full waterproofs). Buchak suggests that blame is different – it's an on-off reaction, rather than something we increase or decrease along with the evidence. For example, suppose you know that one of your two children created a huge mess but, until you interrogate them – and see their guilty or indignant faces – you don't know which one it was. Suppose they are both scamps so it's 50 per cent likely

²⁰ I discuss this in [Ross 2023a](#).

²¹ I discuss various ways that this idea has been developed in [Section 4](#).

²² [Buchak 2014](#). Also, see [Littlejohn 2020](#).

either way. The rational thing to do isn't to blame them *both* to a 0.5 degree! Rather, you wait until you have the evidence that allows you to fully believe one was responsible. Buchak concludes that merely having a probabilistic estimate about someone's culpability is not the sort of attitude that justifies blame – *believing that they did it* is the attitude needed for blame.

Criminal courts blame people for breaking the norms of the criminal law. If there's a close connection between blaming and belief in individuals, this might be helpful in our search for a way to understand the standard of criminal proof. Does the belief–blame connection straightforwardly show that the standard of proof needs to be high enough to make sure that the evidence makes it rational to fully believe the accused is guilty?

This is a tempting line of thought, but it's too quick. After all, courts are not the same as individuals. And legal verdicts are not beliefs. Just because an individual might need to fully believe something in order to blame someone, courts are different from human minds. While the standard for blame in an individual mind might be hard (or impossible) to change, we can change the standard of proof at will – a legislator can set the standard of criminal proof at any level they like. So, why would facts about individual blame constrain the courtroom standards of proof?

The answer, I think, returns to the social role of the criminal court. Imagine that criminal courts routinely found people guilty on the basis of evidence too weak to persuade people in the community to believe that the person was guilty. If this was to become common knowledge, there would be a huge tension between the legal system blaming and punishing following a guilty verdict, and the fact that people in society wouldn't personally feel comfortable blaming the accused. This would be a recipe for the criminal justice system to become alienated from the community that it is supposed to represent; courts would no longer be viewed as holding people to account on behalf of the community.

If courts did not aim to bring people to believe in the guilt of the accused – and potentially to blame them – criminal justice would be more like a system of risk management than a moral practice. Sometimes it is acceptable to treat people just as a vector of risk, especially in emergency situations. Suppose there was a highly contagious and fatal disease. In such cases, it might be acceptable to force people into quarantine even if there was only a 50 per cent chance they are infected. But criminal justice, in my view, should not simply be a way of managing risk. A community punishing someone for committing a crime is not analogous to forcing them into quarantine because they *might* have a virus. Rather, criminal justice inescapably involves moral ideas of blame and criticism, where we hold people responsible for falling short in their conduct. If this

is right, we need to make sure that the community can get behind the moral judgements made by criminal courts. Otherwise, we would be taking away people's liberty without members of the community being confident enough to judge them blameworthy, even though the very idea of the punishment is predicated on the person *being* blameworthy. Indeed, without belief in the blameworthiness of the accused, it would be hard for a community to view punishment as legitimate. To maintain the apparent connection between moral blame and criminal conviction, we require a standard of criminal proof strong enough to support a community-wide belief in the guilt of the accused. This means that the standard needs to be rather demanding.

That is my argument for a demanding standard of proof, irrespective of Laudan's claims about the harm of reoffenders. Of course, saying that the standard must be strong enough to support belief in the guilt of the accused leaves many questions open. For example, I have not shown that there is any equivalency between 'beyond reasonable doubt' and the standards for rational belief – believing something may not be the same as believing it beyond reasonable doubt.²³ It seems quite plausible that rational belief is necessary for believing beyond reasonable doubt, but not sufficient. We'll get to some of these questions in [Section 2](#). But what we have is a good start, a lower limit on how weak a standard of criminal proof can be before it begins to undermine the very purpose of criminal justice.

1.9 Back to Civil Proof

What about the civil standard of proof? Although I mainly focus on criminal law, proof in civil law is just as socially important and philosophically difficult. One big difference between criminal and civil law is that civil law doesn't necessarily involve the severe punishments of the criminal law. But for most people, their lives and relationships (personal, economic, social) are structured more fundamentally by the wide-ranging rules of civil law than by criminal prohibitions.

The civil standard of proof – 'more likely than not' – strikes a different balance between false positives and false negatives than the criminal standard. Rather than attempting to minimise false positives, the civil standard yields a more even balance of risks. The civil standard of proof could therefore be

²³ It is interesting to note that in England and Wales judges are directed to ask jurors to convict only if they 'are sure' of guilt. This state of 'being sure' is regarded, in theory, as the same as beyond reasonable doubt.

taken as indicating *indifference* between false positives and false negatives.²⁴ Such indifference might indicate, for instance, that the court views it equally important: (i) to avoid mistakenly holding employers liable for negligence as (ii) to avoid mistakenly overlooking employees who are harmed by negligent employers.

There are many cases where indifference between false positives and negatives in civil law seems sensible. For instance, if two suburban neighbours are arguing about whether *A*'s garden hedge is encroaching *B*'s property, we might think that there is no reason to stack the deck in favour of either party. It might be inappropriate for the state to be seen taking sides in a dispute between private parties by – for example – forcing one side to prove their point beyond reasonable doubt.

But in many civil cases, the justifiability of indifference between different mistakes is less clear. Not all cases are between private individuals and often there are inequalities of power that we might want the state to care about. For example, the same standard of proof is used in disputes about garden hedges as in cases involving the removal of vulnerable children from the family home. Perhaps the badness of exposing a child to abuse is roughly equal to the badness of unnecessarily taking a child away from its parents. But this is not obvious.²⁵ And in other cases, we might think that some mistakes are more harmful than others. For example, in a civil case involving protection of the environment, we might think that failures to identify mass pollution are especially harmful compared to the cost of unnecessarily making a corporation improve its environmental protection practices.

Civil law has a more ambiguous relationship to moral blame than criminal law. Some civil cases do involve holding people responsible for conduct we would ordinarily regard as morally blameworthy. For example, many serious crimes – such as sexual crimes – can also be pursued in civil courts. Moreover, some civil cases, especially in the US, can lead to ‘punitive damages’ where the losing party does not simply compensate the other side for their estimated loss, but pays an excess as a quasi-punishment. This raises puzzling questions – if we think that blame is the thing that makes the criminal standard so high, why do the same arguments not apply to the civil law? Of course, the fact that current legal practice does not fully fit with our best theories is not always a reason to think the theories are wrong. Legal systems evolve in a piecemeal and sometimes haphazard way; we should not be surprised by the existence of awkward

²⁴ This itself is not clear. Given that *more than one thing* generally needs to be proven to win a civil case, the burden of establishing multiple points on a given standard of proof might make it harder for the party bringing the claim to win.

²⁵ See *Re. B* [2016] UKSC 4 for discussion in case law.

cases that do not fit with our general understanding. But the philosophy of civil proof is an area where there is a great deal of important philosophical work still to be done.

2 Proof: Fixed or Flexible?

The standards of proof appear to be surprisingly inflexible. Take criminal law. The very same standard is used to determine guilt for war crimes in the International Criminal Court and for petty theft in your local court. In each case, the standard is proof 'beyond reasonable doubt'. There is a similar inflexibility in many areas of civil law. The same standard is used to decide whether your garden hedge is overgrown and whether the state should remove a child at risk from their family home. In each case, proof is 'on the balance of probabilities'.

From one perspective, such inflexibility might seem fair and reasonable – everyone is treated equally in court, putative genocidal maniacs and petty thieves alike. But from another perspective, it is strange. Generally, when making decisions in our everyday lives, we change our approach depending on the type of decision we are making. You will want to be more certain when making a bet involving a year's salary compared to a £20 flutter on Saturday's football scores, more certain about weatherproofing when buying a house compared to renting an Airbnb.

Just like decisions in everyday life, decisions faced by criminal and civil courts vary in their gravity. It might be hard to specify all the ways in which cases can differ, but it is uncontroversial to say that some cases involve bigger risks than others. For example, a finding of guilt for petty theft might cause the accused to receive a small fine, while for murder it might lead to them receiving a long prison sentence or even the death penalty. This raises the question: should the legal system use different standards for different cases? We will focus on criminal law primarily. So, we begin by asking: should criminal justice systems use different standards for different crimes?

2.1 Different Standards for Different Crimes?

The idea that some crimes require special treatment has a long history. Consider the crime of treason, which historically came with liability to especially gruesome punishments – most infamously, being hung, drawn, and quartered. For long periods in the history of English law, you could only be convicted of treason through the testimony of *two* eyewitnesses, a rule not applied to other crimes. Curiously, this remains codified in the US Constitution – the drafting of